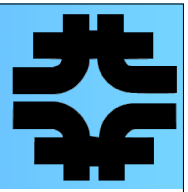

US CMS Goals and Milestones for 2004

LATBauerdick, Fermilab

Joint Steering Meeting
Univ. of Chicago, May 20, 2004

DC04 successfully finished

The CMS 2004 Data Challenge is over -- what's next?

- "The purpose of this milestone is to demonstrate the validity of the software baseline to be used for the Physics TDR and in the preparation of the Computing TDR."
  "[it] comprises completion of a "5% data challenge", which successfully copes with a sustained data-taking rate equivalent to 25Hz ..."
  The emphasis of the challenge is on the validation of the deployed grid model on a sufficient number of Tier-0, Tier-1, and Tier-2 sites."

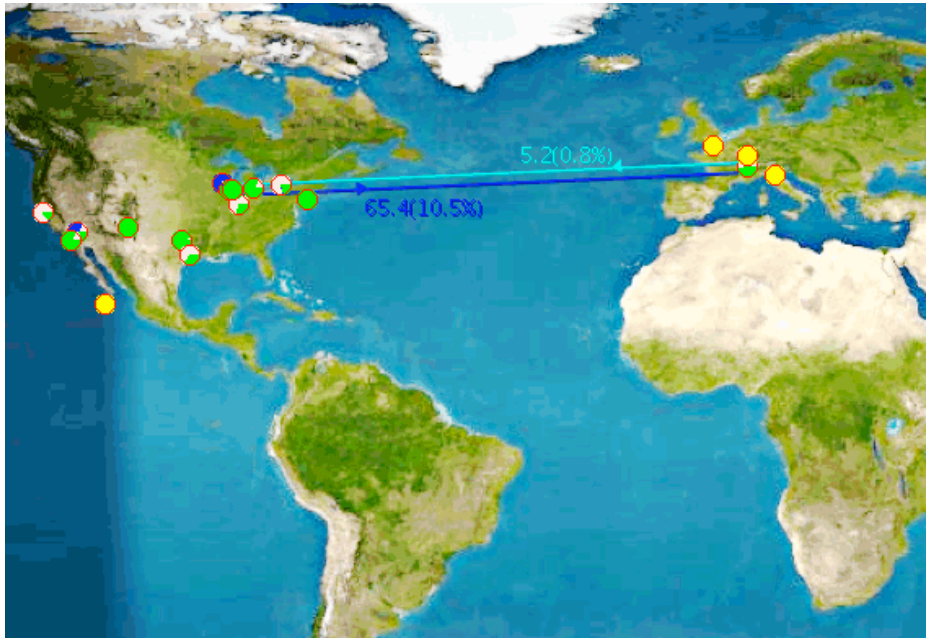

After completion DC04 milestone, 3 years left until CMS t0

DC04: managing many files

A glimpse on what 2007 will be about -- but rationalize data model

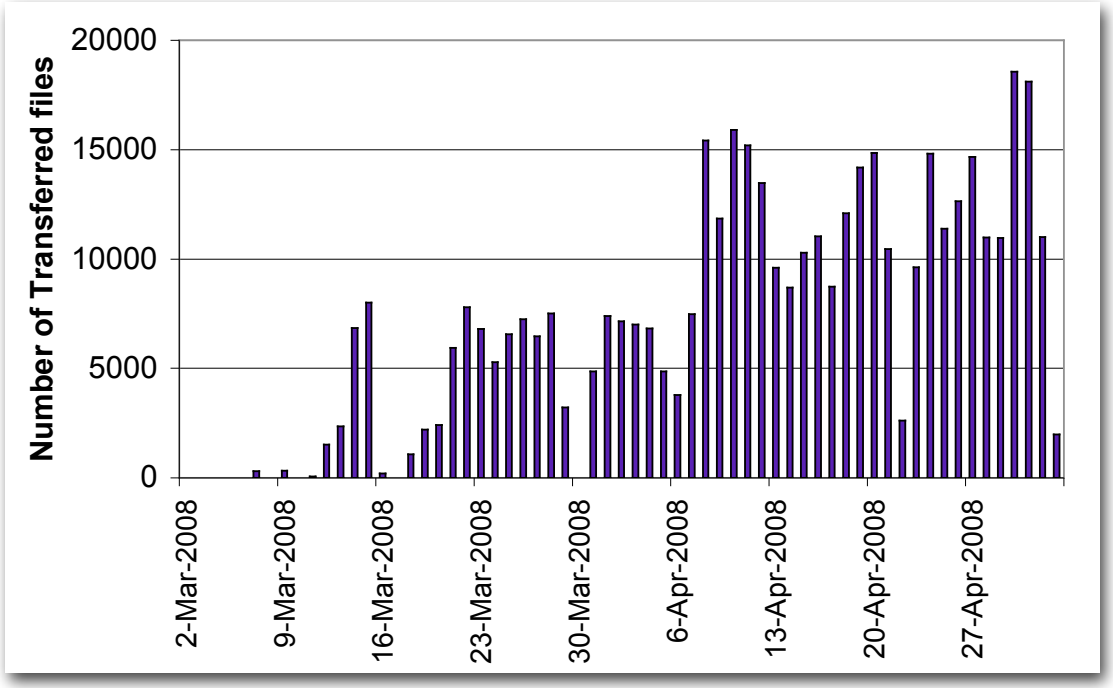

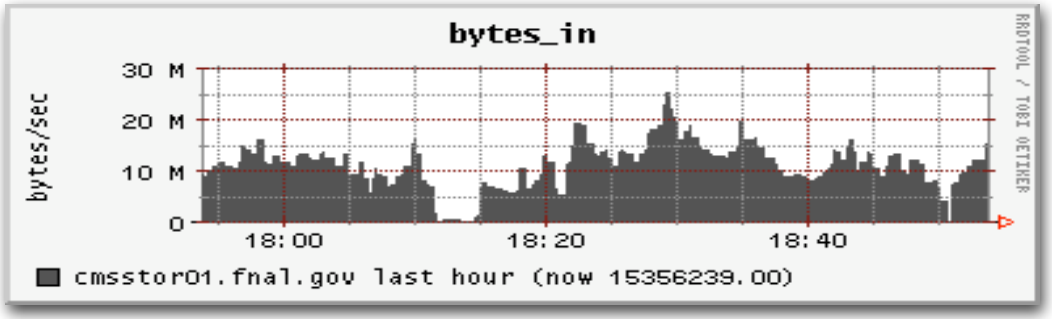

After DC04: New Focus on Data

DC04: production and data "streaming

- largely scheduling of pre-defined jobs across Grid3
- point-to-point data transfers
- now shift to serving data to physics groups for data analysis

data handling, data management, data validation, data access

- consolidate and validate CMS data sets on Grid (Tier 1/2)
- provide and optimize access to large data sets for data analysis
- provide data set discovery through meta-data queries

organize and re-factor resources into "services" available to users

- that together provide end-to-end system for analysis

GAE services for analysis -- emphasis on interactive analysis (Caltech et al.)

experimenting with Virtual Data (U.Florida et al.)

facility and fabric services:storage, data transfers, VO mgmt (Fermilab et al.)

US CMS priority

setup an end-to-end system with the goal to

- make data available to CMS physicists for "offline analysis"
- address data handling, data management, data access issues

provide and further develop the required fabric resources

- farms and compute services, including "opportunistic" resources
- disks/tapes and storage services, data movers and data access
- VO management and AAA services, security and policy infrastructure

application level services: handling / management of CMS data sets and jobs

- data set services
- catalogs: file metadata, replica, event and data set meta data
- distributed file delivery and caching services
- storage services: permanent and transient, managed and reliable
- management and tracking of processing/analysis jobs

learn from previous experiences

CMS-ARDA and EGEE/VDT middleware

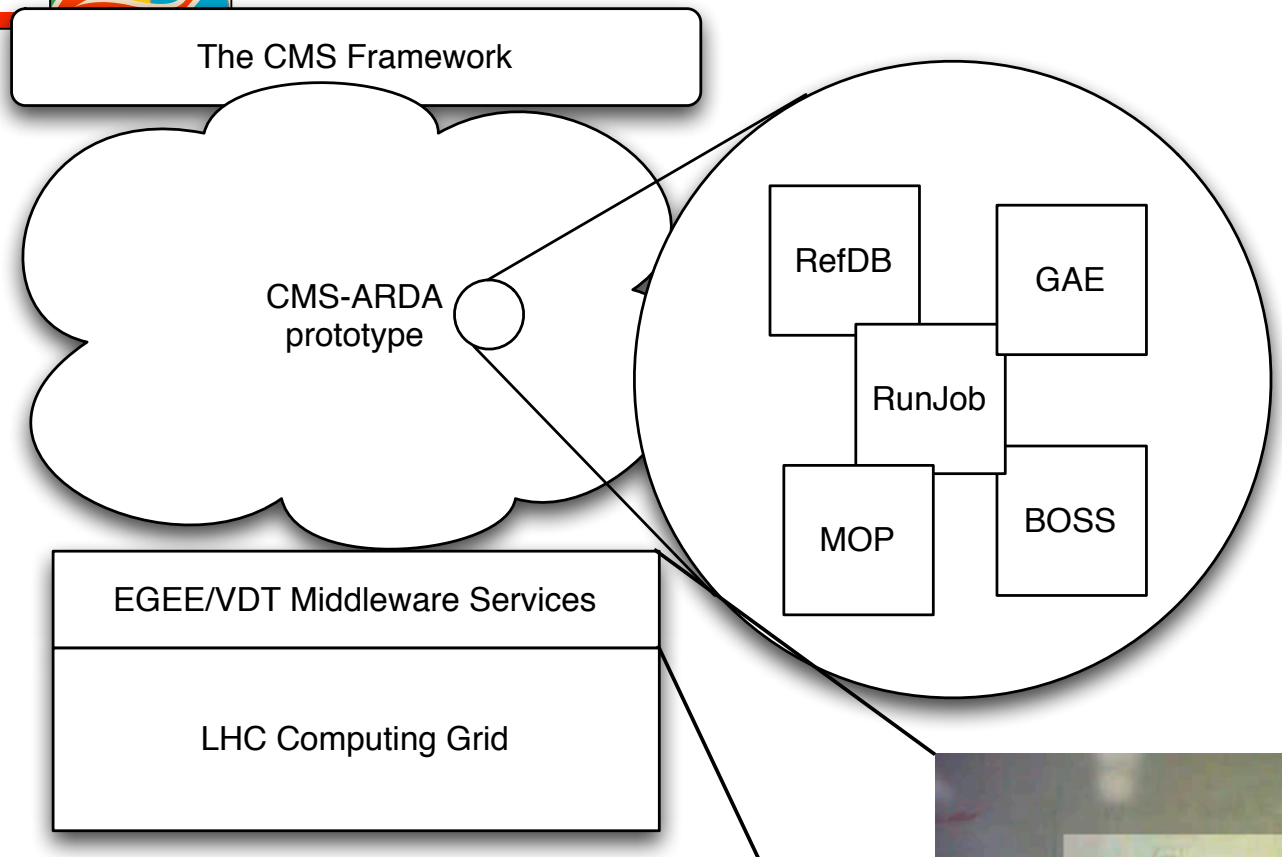

Application-level services interfacing to COBRA and ARDA/ EGGE MW

CMS to diagonalize into engineered end-to-end system

"prototype of generic middleware stack capable of supporting prototypes of distributed analysis environments of experiments, powerful enough to support end-to-end capabilities required"

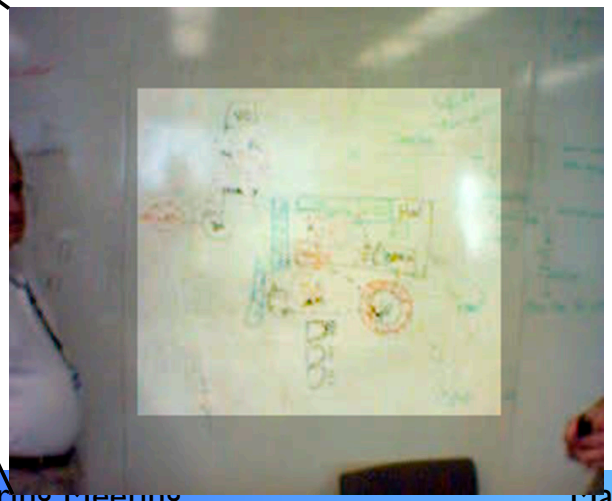

End-to-end Data Handling System

CMS will need to have end-to-end data handling, and soon

to build a system, we better become rather clear on the requirements

that needs baseline decisions on computing and event data model

- e.g. sequential access to files? data tiers? "object pointers" across files (maybe zipped files) the rule or the exception? does the AOD replicate the ESD objects as needed or ref to it? the ESD/DST replicate the MCinfo or ref to it? Etc
- e.g. event-level and data-set level metadata part of "external catalogs" like POOL user collections? Or are components like conditions database framework, POOL relational database services, Grid replica catalogs and services, etc to be integrated into "the CMS data handling system"? Or the other way round?

The system will need considerable engineering, and we need to agree on an approach to how to run this in CMS very soon

"Globally Enabled Local Data Analysis"

Approach: *Think Global Act Local*

USCMS approach to enable local use,
with a view to extend into a truly distributed system

CMS proposed Nov2004 milestone looks consistent with that view:

- "ability of a user logging in to a "local" T1/T2
- to locate the whereabouts of data-sets of interest to him
- and to submit via a GRID interface (not by remote explicit login)
- two types of task: jobs to process the data remotely
  - and publish the results,
- and jobs to copy data from a remote T1/T2 to another T1/T2
  - and run the analysis jobs at this destination center."

We are starting to determine how to break the use case down
into a set of application services services to implement it
based on the USGrid/LCG/EGEE/GAE/ARDA environment

Building up the U.S. LHC Fabric

Expect to start deployment of ~5 Tier-2 centers this year

Challenges with funding -- opportunistic use for production

Open Science Grid strategy

Interoperability milestone with LCG imminent:
job scheduling, data management, VO management

close integration much harder: packaging, AAAA, policies,
VDT/EGEE midleware services

Transferring LHC Data to the U.S. Grid

Streaming and On-demand bulk transfers:

- Storage Interfaces, SRM, Data Transfers Network services, "LambdaStation" data model, physics selection

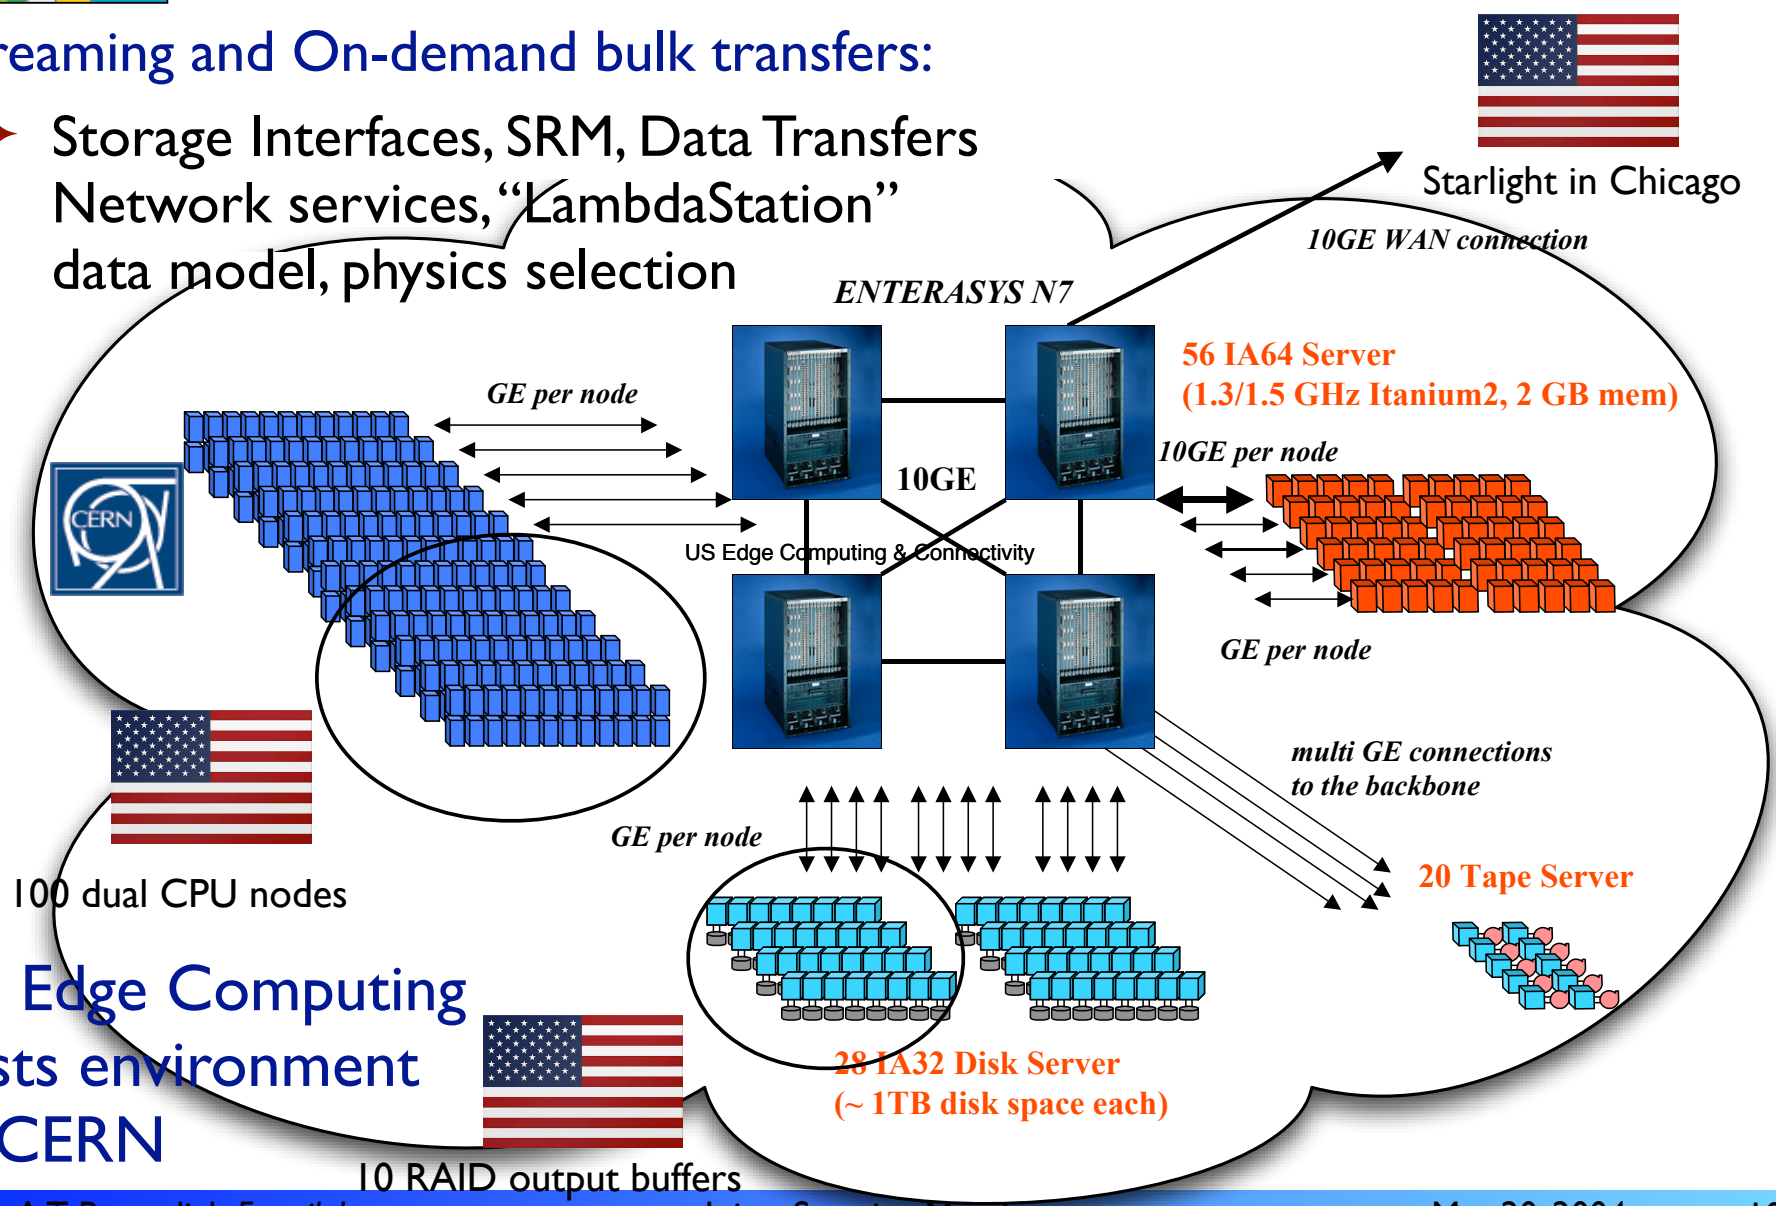

US Edge Computing Tests environment at CERN

Concrete E2E Demonstrators

Production 10Mevents/Month world-wide

- end-to-end multi-grid production
- application level data management
- inter-grid job scheduling

Analysis Demonstrators

- end-to-end prototypes
- data services, GAE, CMS-ARDA/EGEE/VDT, Virtual Data

US Edge Computing

- end-to-end data transfer

For all these we rely on the help of Trillium

Alignment with the Open Science Grid Roadmap